

# Observations from the Deployment of a Large Scale PKI

Rebecca Nielsen  
Booz Allen Hamilton

## 1 Background

The United States Department of Defense (DoD) has been investigating the use of public key technology to help meet its information assurance goals since 1997. The DoD implemented a pilot Public Key Infrastructure (PKI) in 1998, and began a mass rollout of the current DoD PKI in 2000. Since then, the DoD has successfully issued digital certificates on Common Access Cards (CAC) to over 85% of its 3.5 million user population. While the deployment of the DoD PKI has not always been smooth, the issuance of digital certificates has been one of the first truly enterprise-wide standard technology implementations within the DoD.

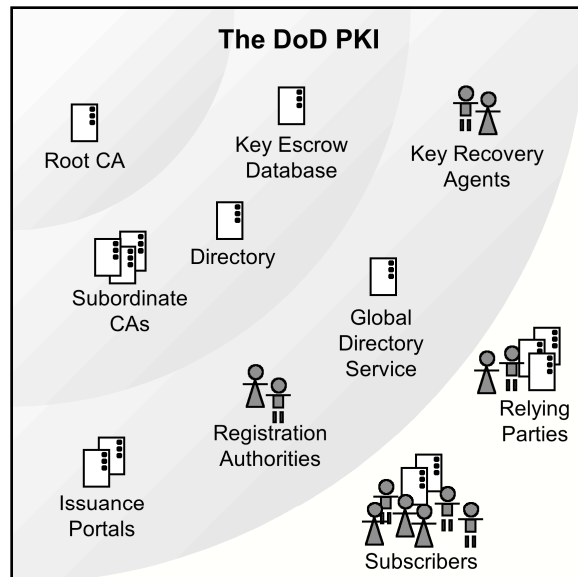
This paper provides insight into some of the technology and organizational lessons learned in deploying the world's largest PKI from the perspective of a DoD contractor.

## 2 Managing the "I" in PKI

### 2.1 The DoD PKI Architecture

The DoD PKI consists of a single Root Certification Authority (CA) and multiple subordinate CAs. The Root CA only issues subordinate CA certificates.

Subordinate CAs issue five types of certificates: identity, signature, encryption, component, and code signing. Identity and signature certificates can be used to authenticate to applications or digitally sign forms or email messages. Because many DoD email addresses change when individuals move from one location to another, the DoD issues the primary identity certificate with no email address. The signature and encryption certificates contain email addresses to support Secure/Multipurpose Internet Mail Extensions (S/MIME) version 2 and 3. New email certificates can be issued based on the presentation of



a valid identity certificate. Component certificates are issued to web servers and other devices. Code signing certificates are issued to specific entities within the DoD that approve mobile code.

Other DoD PKI core components include the internal directory servers and the Key Escrow Database (KED). All private keys associated with encryption certificates are escrowed prior to the issuance of the certificate.

The DoD PKI supports two interfaces for certificate issuance, hardware and software. Hardware certificates are issued on the CACs to all DoD personnel. The CAC is a Java smart card that has been validated against Federal Information Processing Standard (FIPS) 401 Level 2 requirements. The Java card was selected both to allow for the inclusion of additional functionality beyond PKI on the card and to enable multiple vendors to provide card stock to the DoD. Identity proofing and certificate issuance on the CAC take place using the DoD's existing personnel identification card issuance system. Since the CAs do not have a direct method for interfacing with the CAC, issuance portals are used to facilitate key generation, certificate request generation, and insertion of issued certificates.

Software certificates can be issued to people and web servers. Although the CAC is the primary issuance process for personnel, software certificates are used

to support some legacy applications that do not yet support hardware tokens, and in some environments where CAC issuance is difficult. Software certificates are requested via a Hyper Text Transfer Protocol, Secure (HTTPS) interface and verified by Registration Authorities (RA) and Local Registration Authorities (LRA).

For publication of PKI information, the DoD PKI interfaces with the Global Directory Service (GDS). GDS is an internal enterprise directory that supports both HTTPS and

Lightweight Directory Access Protocol (LDAP) interfaces. Subordinate CA certificates, Certificate Revocation Lists (CRL), and encryption certificates are published to the GDS from the DoD PKI. Subordinate CA certificates and CRLs are also

published to an external X.500 directory for access outside of the DoD.

Certificate revocation is performed by RAs using an HTTPS interface. Key recovery is performed by key recovery agents who access the KED using an HTTPS interface.

## 2.2 Certification Authority (CA) Scalability

Take into account all of the required tasks when developing architecture requirements.

When designing the infrastructure, the DoD performed load testing to determine how many certificates a given CA could issue. However, this initial load testing did not account for the many other functions that CAs must be able to perform at the same time as certificate issuance, including the following:

- validating credentials of trusted personnel,
- publishing certificates,
- generating CRLs,
- publishing CRLs,
- revoking certificates,
- responding to requests to search for specific certificates.

The DoD PKI has over 2,000 approved RAs. Because of personnel turnover, the list of approved RAs changes frequently. To ensure that certificates can still be issued if one or more CAs are unavailable, the DoD PKI is configured so that all RAs are authorized on all CAs. The interface provided by the vendor to manage trusted personnel did not support the large number of RAs or the requirement for frequent updates. To minimize the impact to CAs, the DoD has developed custom scripts that allow changes in RA personnel to be quickly uploaded to all CAs.

The process of generating CRLs requires significant CA processing time, both in determining which certificates have been revoked, and, as the CA ages, determining which revoked certificates have expired and should not be placed on subsequent CRLs. For CAs that issue a large number of certificates, the requirement to check each revoked certificate for its expiration date can cause the total time to generate a CRL to be greater than the next update period of the CRL. While CRL generation requires less processing time if expiration date checks are not performed, continuing to include expired certificates on CRLs increases the overall CRL size.

Although the GDS is the primary interface for applications to retrieve CRLs and for end users to search for encryption certificates, direct searches of the CA internal databases are still required, primarily for certificate revocation. When requesting certificate revocation, most users and supervisors do not know the CA and serial number for the certificate that needs to be revoked. As a result, the RA must search multiple CAs to locate the correct certificate prior to authorizing its revocation.

## 2.3 Hardware and Software Maintenance

PKI cycle times are significantly different than hardware and software cycle times.

### 2.3.1 CA Hardware and Software

The DoD PKI was designed for the long term. The DoD Root CA has a validity period of thirty-six years. Each subordinate CA has a validity period of six years. Subordinate CAs issue certificates for the first three years of their validity period and are then “retired” so that they only issue CRLs for the remaining three years. CAs are only taken completely out of service once all certificates issued by the CA have expired. CAs issued to Government personnel are valid for three years, while those issued to contractors are valid for up to one year.

In contrast, hardware life cycles are one to three years, and software product cycles can be eighteen months or less. As a result, neither the software nor the hardware in use for the DoD Root CA are still supported by their respective vendors. Older subordinate CAs are also operating on non-supported versions of hardware and software, increasing both the requirement for and the cost of maintenance.

### 2.3.2 Key Length

In addition to product life cycles, the basic technology behind PKI is also changing. When the Root CA was established, 512-bit keys were still in use, and 1024-bit keys were the longest supported by vendors. Today, 1024-bit keys are standard, and the Federal Government has published guidelines that all certificates that will expire later than 2008 should be issued with 2048-bit keys.

As a result, the DoD PKI is currently working on a solution to upgrade the Root CA. There are two options for upgrading, migrating the current Root CA to newer hardware and software versions, or establishing a new Root CA and issuing a rollover certificate from the current Root CA to the new Root CA. Migrating the existing Root CA is simpler for

the short term, but does not solve the problem of the 1024-bit Root CA signing key length.

Establishing a new Root CA with a 2048-bit signing key will require pushing down the new key to all applications relying on certificates from the DoD PKI. In addition, a new Root CA will require maintaining two infrastructures for three to six years, depending on whether all current subordinate CAs are retired when the new Root CA is established.

### 2.3.3 Certificate Profile

Another issue with the long term nature of the PKI is that changes to certificate profiles require over three years to implement. For example, the DoD PKI initially did not support the extensions required to use digital certificates to authenticate to Microsoft Windows-based networks. Windows requires the following extensions in certificates<sup>1</sup>:

- CRL Distribution Point must be present,
- Key Usage must be set to Digital Signature,
- Extended Key Usage (EKU) must contain the Smart Card Logon Object Identifier (note that if the EKU extension is populated, it must also contain the identifiers for all other uses for the certificates, such as Client Authentication),
- Subject Alternative Name must contain a User Principal Name (UPN) of the format user@name.com.

Once the requirements were determined and the changes implemented at the subordinate CAs, all new CACs contained a signature certificate with the additional information. However, there are still some subscribers within the DoD that do not have the required extensions on their CACs.

Another example is the Authority Information Access extension. Research by the Federal Bridge Path Discovery and Validation Working Group has indicated that path discovery is facilitated when certificates contain the Authority Information Access (AIA) extension<sup>2</sup>. The DoD PKI does not currently include the AIA extension. If the DoD makes a decision to modify its certificate profiles to include the AIA extension, the change will take three years to be reflected in all DoD PKI issued certificates.

### 2.3.4 Smart Card Technology

In addition to CA and certificate profile updates, the DoD must manage user smart card migration issues. Since most CACs are valid for three years, CAC middleware must concurrently support three years of smart card technology. The DoD is investigating upgrading new card stock to a 64k chip, instead of

the 32k chip currently supported. This new chip will support additional capabilities beyond PKI. The additional space may also support better security protections, which would enable users to perform more card maintenance, such as certificate update, from their own workstations instead of having to return to an issuance station. However, all DoD users will not be able to take advantage of these new capabilities until all cards have been replaced through normal expiration.

## 2.4 Personnel

Integrating PKI rollout with existing processes is a requirement for success.

The initial DoD PKI rollout was planned as a software-based implementation. The DoD would centrally manage the PKI core components, and each DoD service or agency would provide personnel to act as RAs and LRAs who would register individuals. When early adopter applications tried to get their users registered to get certificates, however, they found that RAs and LRAs were not available. Local commands resisted the requirement for additional personnel, and the travel costs for sending RAs and LRAs to training.

At the same time the PKI was performing initial rollout, the personnel office was developing a new identification (ID) card to be rolled out to all DoD military and civilian employees. In November 1999, the DoD made a decision to combine the two programs and use the new ID card as a hardware token for digital certificates. As a result of this decision, the PKI and personnel offices worked together to design a process that used existing ID card issuance stations to verify identity and issue certificates in conjunction with ID cards.

This new process did increase the personnel requirements for ID card issuance stations. Prior to the CAC, DoD ID cards were only issued to military personnel, but CACs are issued to military personnel, civilian employees, and on-site contractors. Also, the time to issue CACs is longer than the time to issue the old ID cards. However, combining certificate issuance with ID card issuance allowed the PKI to take advantage of the existing ID card infrastructure and minimized the personnel requirements for services and agencies. Also, the requirement for personnel to get a new ID card facilitated the issuance of certificates.

### 3 Technology Challenges

Building the capability to support the DoD enterprise is not sufficient for the success of the DoD PKI. Since PKI is an infrastructure technology, it does not solve any operational requirements unless public key technology is integrated into applications. This section explores the two most significant challenges that the DoD PKI has experienced in gaining acceptance from the functional community for the use of PKI.

#### 3.1 Certificate Status Checking

Checking certificate revocation status is the most difficult technical challenge of PKI.

CRLs are theoretically elegant. They provide a mechanism for a CA to state that it no longer asserts the binding between the identity in the certificate and the associated key pair for a set of certificates that it issued. CRLs provide only a minimum set of data, the certificate serial number, the date of revocation, and optionally a reason for revocation. Because they are digitally signed, the transmission mechanism does not itself have to be trusted in order to accept the information contained in the CRL.

In practice, however, relying on CRLs has not worked well. Products that are enabled to use digital certificates only provide minimal support for CRLs. In some cases, no provision is provided to automate the downloading of CRLs. Vendors who do provide an automated update capability may not allow setting when the attempt to retrieve a new CRL occurs, which can result in multiple applications attempting to access the CRL repository at the same time. At least one vendor treats the next update field of a CRL as an expiration date, and will not validate any certificate issued by a CA for which a current CRL is not available. Finally, the information contained in a CRL is only as current as the time the CRL was published, which results in significant latency issues.

The scale of the DoD PKI results in an additional problem, the overall size of CRLs. The combined size of the CRLs from all of the DoD PKI CAs is approaching 40 megabytes. It is not feasible for every application on DoD networks to download this amount of data every day without having a significant impact on available bandwidth.

Although CRLs are an efficient way of publishing revocation information for the entire PKI, no single application has all subscribers as users, so each application only needs a subset of the information. However, the enterprise PKI does not know in advance which subset is needed by each application.

The DoD PKI has examined two alternate CRL approaches, partitioned CRLs and delta CRLs. A CA creating partitioned CRLs divides certificates into blocks of a preset size based on information contained in the certificate such as the certificate serial number. Instead of issuing one CRL, the CA issues multiple CRLs, one for each preset block of certificates. When an application attempts to validate a certificate, it checks to see if a current CRL for the block the certificate is contained in is locally cached, and downloads the CRL partition if it is not. While partitioned CRLs allow applications to only retrieve limited CRL information, the DoD has not developed a solution involving partitioned CRLs, partially because of the lack of support from either CA or application vendors.

A CA supporting delta CRLs issues a full CRL once or periodically, and then only issues delta CRLs that contain certificates that have been revoked since the last delta CRL was issued. As a result, delta CRLs are significantly smaller than full CRLs. However, applications must have a mechanism of ensuring that they have downloaded all delta CRLs, because no single CRL can be considered an authoritative source for information on the revocation status of any given certificate. Although the DoD PKI does not support delta CRLs, one agency within the DoD has successfully piloted a delta CRL approach for transmitting revocation information in a severely bandwidth constrained environment.

In general, CRLs have been an effective method of transmitting revocation information across enterprise networks with high bandwidth availability, but are too cumbersome to use to get this information down to individual applications.

As a result of the continued issues with performing certificate validation using CRLs, the DoD PKI is deploying an infrastructure to support revocation status checking using the On-line Certificate Status Protocol (OCSP). To meet the requirements for decentralization, availability, and scalability, this infrastructure will not interface directly with the DoD PKI CAs. Instead, it will provide a capability to download CRLs and provide real-time OCSP responses from multiple locations across the DoD network. Instead of downloading CRLs, applications that support OCSP will be able to get real-time responses for specific certificates from this global robust certificate validation system. Although the use of CRLs as the authoritative source for revocation information does not address the latency issues of CRLs, this hybrid approach of CRLs and OCSP will take advantage of the efficiency of CRLs

and provide an interface for applications that is easier to implement and maintain.

### 3.2 Key Recovery

The person most likely to need key recovery capability is the subscriber.

The DoD PKI key escrow and recovery solution was designed when the DoD PKI was primarily issuing software certificates. The solution was designed to support situations when a private key needed to be recovered by a manager or law enforcement agent to access encrypted data, and occasionally by subscribers who had lost their private keys. As a result, the process was manual and personnel intensive, requiring first that requestors verify their own identities and provide justification for the request, and then that two key recovery agents together retrieve the escrowed keys out of storage.

The transition to the CAC as the token for key generation and storage created a significant change in the key recovery requirement. Since the private key associated with each subscriber's encryption certificate is stored only on the CAC, subscribers lose access to their own private keys at CAC expiration because the old CAC is surrendered at the time of new CAC issuance. In order to access files previously encrypted, all subscribers must recover their old private keys.

The manual process which was designed to prevent abuse of key recovery is too costly to support for a large number of individuals requesting their own escrowed private keys. Since individuals are assumed to be authorized to have access to their own keys, a more automated process is being developed that will allow subscribers to use their identity certificate to authenticate to the key escrow system and request retrieval of their own private keys.

## 4 Organizational Challenges

Although the implementation of the DoD PKI has experienced technical challenges, overcoming organizational obstacles has sometimes proved a harder task. Some of these obstacles are independent of PKI, such as issues relating to coordinating common processes across the worldwide enterprise, developing working relationships between disparate commands within the enterprise, and determining which organizations within the enterprise should have primary responsibility for each element of the overall architecture. This section highlights some of the organizational challenges specific to PKI implementation for users, managers, and developers.

### 4.1 The Users

Provide users with new capabilities that help them to get their jobs done, not just PKI certificates.

Most users will embrace new technology if they see a clear benefit to its use. However, PKI was initially marketed as a technology, not as a mechanism for getting the job done. For example, "PKI 101" training often starts by stating the concepts of public key technology, then introduces the user to "Alice" and "Bob" who are exchanging signed and encrypted email messages. By this time, attendees have decided that PKI is very complicated and since they can send email without PKI (and have been doing it for years without problems), they leave the training having decided that PKI is too difficult.

The DoD PKI was originally targeted as a pilot that would provide better assurance for a new electronic travel system. Through the use of digital signatures, travel claims could be processed significantly faster, resulting in shorter times for employee reimbursements. Once deployed, the PKI could then be used with other systems. However, delays in the rollout of the travel system and the decision to implement a smart card-based PKI meant that many users were issued a CAC months prior to receiving a smart card reader and without any application requiring use of their new certificates. As a result, most users' experience with PKI consisted of waiting in line to get a CAC and then using the CAC the same way they had used the ID card they had prior to the CAC. These users did not see any real benefit to the new technology.

Although smart card readers are being deployed and applications are beginning to incorporate support for public key technology, the DoD PKI continues to struggle to attain widespread user acceptance.

### 4.2 The Managers

Application owners need policy, budget guidance, and a business justification for adopting public key technology.

Within the DoD, funding for PKI core components and card stock is centrally managed. However, individual applications, including email, networks, and web servers, are very decentralized. Therefore, rolling out the PKI required a few decisions by policy makers, but integrating public key technology into applications requires many decisions by many application owners. These managers must consider

multiple demands when determining how to allocate limited resources:

Getting support from application managers requires providing managers with the information they need to make decisions including the following:

- Ensure that published policy is consistent with the organization's goals for integrating public key technology. Policy enables early adopters of new technology to justify their investments.
- Provide direction for requesting funding for public key enabling as a part of the standard budget cycle. Getting out of cycle funding to meet security driven requirements is difficult and almost always means that other planned functionality must be sacrificed to meet PKI requirements.
- Use specific examples when presenting security requirements to application owners to show what vulnerabilities exist in current systems and how the use of PKI can help to mitigate them.
- Define business case benefits for PKI in addition to the "better security" case. Because PKI acceptance has been primarily in the security community, explanations for why to integrate public key technology tend to be heavily security focused. The business case for PKI, including more efficient user management, decreased password management, and new functionality capabilities, should be more clearly stated.

Getting the acceptance of application owners is a critical step in showing a return on investment in PKI. However, most application owners will not commit to enabling existing applications until the users have digital certificates. DoD applications that made initial investments in using PKI in the late 1990s all delayed public key enabling because of the inability of their internal DoD users to register for certificates. Now that the DoD has invested resources to issue certificates to eligible users, these and other applications are finally beginning the transition to using public key technology.

#### 4.3 The Developers

Better training is needed to assist developers in public key enabling.

Ultimately, the success of PKI is dependent on developers performing system integration to public key enable applications. DoD applications usually involve some components, such as web servers, that have native support for some PKI capabilities and other components that do not support PKI. Public key enabling, therefore, can require upgrading software to later versions that support required

capabilities, replacing components with similar components from different vendors that support required capabilities, and building interfaces between enabled components and those that do not support public key technology.

Unfortunately, there are relatively few individuals who understand both PKI and application architectures. Available PKI training consists primarily of lessons on how to stand up the infrastructure; it does not focus on how to integrate PKI into existing applications. Vendors provide some guidance, but this information is usually limited to the vendor's own products.

For example, a web server vendor will provide instructions on how to request and install a server certificate and how to turn on client certificate-based authentication. The vendor may also provide instructions on how to perform certificate validation. However, nothing is provided on how to integrate the authenticated identity information from the certificate with access control to a database or other back end component.

Training targeted at developers on how to integrate PKI into real-world applications should become much more widely available.

## 5 Conclusion

As the DoD has rolled out PKI, it has experienced technical challenges. However, PKI has been more successful than many technologies in meeting the scalability demands of the DoD enterprise. By integrating certificate issuance with existing personnel processes, the DoD PKI has been able to perform in-person identity verification to over three million users. Technology challenges have been met using a combination of redundant systems and customized interfaces developed by DoD, contractor, and vendor personnel working together. The only basic building block of PKI that has not scaled successfully is the CRL. However, the DoD is working to overcome this barrier through the use of OCSP.

The DoD PKI rollout has also encountered issues surrounding the enterprise nature of PKI. PKI implementation has required that existing business process problems be resolved prior to the success of PKI.

The more difficult challenges have been in getting acceptance from the user, application owner, and developer communities. A primary reason for this issue is the lack of good training available targeted specifically to the interests of these communities.

The DoD is committed to continuing down the path of PKI deployment and public key enabling of applications. The implementation of PKI is a long term investment, since application owners do not want to commit to using certificates until they believe that their user population has the capability to get certificates. Unfortunately, the return on investment in PKI does not become measurable until applications have started to use the technology. Staying the course has presented challenges, but the potential of public key technology, both in current architectures and in the next generation Net-Centric environment, is critical to meeting the DoD's information assurance goals and improving its business processes.

---

<sup>1</sup> "Guidelines for Enabling Smart Card Logon with Third-Party Certification Authorities," Microsoft Knowledge Base Article 281245.  
<http://support.microsoft.com/default.aspx?scid=kb;en-us;281245>

<sup>2</sup> "Functional Requirements for Path Validation Systems," Path Validation Working Group, Draft Version 0.8, March 2004.  
<http://www.cio.gov/fbca/pdvalwg.htm>